Mixed Models Analysis of Urbanization Level on Chlorpyrifos Exposure

Peter P. Egeghy[1], James J. Quackenboss[1], P. Barry Ryan[2]

[1] U.S. EPA, Office of Research and Development, National Exposure Research Laboratory, Human Exposure & Atmospheric Sciences Division, Human Exposure Research Branch, P.O. Box 93478, Las Vegas, NV 89193-3478

[2] Dept. of Environmental and Occupational Health, Rollins School of Public Health, Emory University, 1518 Clifton Road, Atlanta, GA 30322

Abstract

The National Human Exposure Assessment Survey (NHEXAS) pilot studies were conducted from 1995 through 1997 to examine human population exposure to a wide range of environmental contaminants. In the NHEXAS-Maryland study, a longitudinal design was used to repeatedly measure aggregate residential chlorpyrifos exposure in a stratified random sample of 80 individuals. Chlorpyrifos is a semi-volatile insecticide which has been found to persist for weeks following application. Until recently, chlorpyrifos was commonly used for indoor and outdoor treatments by both residents and commercial applicators and also as a termiticide during residential construction. It was also present in a number of agricultural commodities. The effect of urbanization stratum (i.e., characterization of participant's community as urban, suburban, or rural) on chlorpyrifos exposure was investigated using mixed-effects regression models to accommodate longitudinal data and to estimate variance components. Three surrogates of exposure, namely, indoor air concentrations (n = 97), surface dust loading (n = 123), and creatinine-adjusted urinary metabolite (3,5,6-trichloro-2-pyridinol) concentrations (n = 341), were considered separately. The effect of the specified covariance structure was explicitly evaluated to determine if the pooling of variance components, as is often performed by default, is appropriate. Likelihood ratio tests comparing possible covariance structures suggested that the pooling of within- and between-person variance components among urbanization strata may not be appropriate for indoor air concentrations ($p < 0.001$). The variance estimates indicated that indoor air concentrations were far more variable among households in rural communities than in urban or suburban communities. Furthermore, since the estimates of the arithmetic means depend upon the variance components, the choice of covariance structure exerted a large effect upon the estimates of indoor chlorpyrifos concentrations. Mean indoor air concentrations for urban, suburban, and rural households in this study were estimated to be 7.3, 20.2, and 23.0 ng/m³, respectively, when allowing for distinct variance components, but were estimated to be 11.9, 16.6, and 11.7 ng/m³, respectively, when pooling the components. Similar analyses, on the other hand, indicated that variance components may be pooled among urbanization levels when evaluating surface dust and urinary metabolite concentrations. These preliminary results, albeit based on relatively small sample sizes, indicate that the choice of covariance structure can have a large effect on the results of an analysis and must be given appropriate consideration.

Background

In the analysis of longitudinal data there is a need to accommodate the covariation induced by repeated measures on the same subject. With mixed-effects statistical programs such as SAS PROC MIXED, users can select among several covariance structures (for example, compound symmetry, random coefficients, and first-order autoregressive) to model their data. Additionally, the degree of heterogeneity in the covariance structures can also be specified.

The default assumption in PROC MIXED, which is often applied automatically and uncritically, is that both the between-subject and within-subject variance parameters can be pooled across groups. In other words, only two variance parameters, one between-subject component and one within-subject component, are estimated and applied to all subjects, irrespective of grouping level. Two levels of heterogeneity, however, can be specified to accommodate different variance patterns across groups: 1) distinct between-subject but common within-subject variance, or 2) distinct between- and within-subject variance.

Although pooling information across groups increases the precision of variance components estimates and allows analysis where some groups contribute few repeated measurements, important errors can be made when assumptions regarding homogeneity are not met. Accurately modeling the heterogeneity in the covariance is important because valid inferences about the estimated parameters can only be made when an appropriate covariance model is selected.

Qualitative indices of relative goodness-of-fit such as Akaike's information criterion (AIC) and Schwarz's Bayesian criterion (SBC) are readily available and are often used to compare models with different covariance structures. Some (Rappaport et al., 1999; Weaver et al., 2001) have suggested that a more rigorous approach, the Likelihood Ratio Test (LRT), be used. The LRT is straightforward and is easily performed using common spreadsheet programs such as Excel.

Methods

Data from the National Human Exposure Assessment Survey in Maryland (NHEXAS-Maryland) was used. Participants were from three urbanization groups: urban, suburban, or rural. Three measures of chlorpyrifos exposure, chlorpyrifos in indoor air and house dust and the metabolite 3,5,6-trichloro-2-pyridinol (TCPy) in urine, were evaluated to demonstrate the Likelihood Ratio Test and the consequences of inappropriate covariance structures. All dependent variables were found to be lognormally distributed. Group-specific geometric means for each type of measure were estimated by mixed-effects analysis.

Likelihood Ratio Test (LRT):

1. Fit the three models (pooled variance parameters, pooled within-person variance, and distinct within- and between-person parameters) to the log-transformed data using PROC MIXED.
2. Compute the Likelihood Ratio Test statistic (the "Deviance") by subtracting the "-2 Restricted Log Likelihood" value (provided in the "Fit Statistics" section of the SAS output) of one model from another [the difference between two logged values equals the logarithm of the ratio of the values].
3. Determine degrees of freedom by counting the number of variance parameters provided in the "Covariance Parameter Estimates" section of the SAS output for each model and subtracting one from the other.
4. Compare the Likelihood Ratio Test deviance statistic to the Chi-Square distribution.

Choosing a Model:

Unless the LRT indicates that more parameters provide a significantly better fit, choose the model with the fewest parameters, thereby maximizing parsimony and increasing power.

Group Mean and Variance:

The arithmetic mean ($\mu_{x,h}$) and variance ($\sigma^2_{x,h}$) for each of the *h* groups is obtained using the mean of the logged values for each group ($\mu_{y,h}$) from the "Solution for Fixed Effects," the appropriate between- and within-person variance components ($\sigma^2_{B,h}$ and $\sigma^2_{W,h}$), and the following relationships:

Group arithmetic mean:

$$\mu_{x,h} = \exp(\mu_{y,h} + 0.5*(\sigma^2_{B,h} + \sigma^2_{W,h})) \quad \text{and}$$

Group arithmetic variance:

$$\sigma^2_{x,h} = \mu^2_{x,h} (\exp\{\sigma^2_{B,h} + \sigma^2_{W,h}\} - 1)$$

Results

Table 1

Description of the three different covariance structures and the numbers of variance parameters estimated with each model. Note that there are three urbanization groups (urban, suburban, and rural).

Model	Description	No. "between" Parameters	No. "within" Parameters	Total No. Parameters
1	Distinct σ^2_B & σ^2_W	3	3	6
2	Distinct σ^2_B/Pooled σ^2_W	3	1	4
3	Pooled σ^2_B & σ^2_W	1	1	2

Table 2

Mixed-effects estimates of group-specific parameters for three competing covariance structures.

Model	Group	Total Variance ($\sigma^2_{B,h} + \sigma^2_{W,h}$)	Geometric Mean ($\mu_{y,h}$)	Arithmetic Mean ($\mu_{x,h}$)
1	urban	1.4	1.3	7.3
	suburban	2.7	1.7	20.2
	rural	3.9	1.2	23.0
2	urban	1.5	1.3	8.1
	suburban	2.4	1.6	16.2
	rural	4.4	1.2	30.3
3	urban	2.4	1.3	11.9
	suburban	2.4	1.6	16.6
	rural	2.4	1.2	11.7

Table 3a

Likelihood Ratio Tests comparing the three competing covariance models for Chlorpyrifos in indoor air. The model specifying heterogeneous variances (distinct σ^2_B & σ^2_W) provides a significantly better fit than either model pooling the variances across groups.

			1 vs 2:			1 vs 3:			2 vs 3:		
Model	(-)2LL	Parameters	Deviance	df	prob	Deviance	df	prob	Deviance	df	prob
1	294.1	6	18.9	2	0.0001	22.6	4	0.0002	3.7	2	0.1572
2	313.0	4									
3	316.7	2									

Table 3b

Likelihood Ratio Tests for TCPy in urine. With no benefit gained by estimating distinct parameters, both the between- and within-person variance components can be pooled across groups.

			1 vs 2:			1 vs 3:			2 vs 3:		
Model	(-)2LL	Parameters	Deviance	df	prob	Deviance	df	prob	Deviance	df	prob
1	763.1	6	2.6	2	0.2725	7.0	4	0.1359	4.4	2	0.1108
2	765.7	4									
3	770.1	2									

Table 3c

Likelihood Ratio Tests for Chlorpyrifos in surface dust. Despite the large difference in relative group means across the models in Figure 1c, estimating distinct parameters does not provide a better fit.

			1 vs 2:			1 vs 3:			2 vs 3:		
Model	(-)2LL	Parameters	Deviance	df	prob	Deviance	df	prob	Deviance	df	prob
1	440.4	6	2.4	2	0.3012	5.8	4	0.2146	3.4	2	0.1827
2	442.8	4									
3	446.2	2									

Figures 1a-c

Estimated group-specific mean concentrations ($\mu_{x,h}$) of chlorpyrifos in indoor air (a), TCPy in urine (b), and chlorpyrifos in dust (c) from the three covariance models. Note the large effect on group differences resulting from three choice of covariance structure in Figures 1a and 1c.

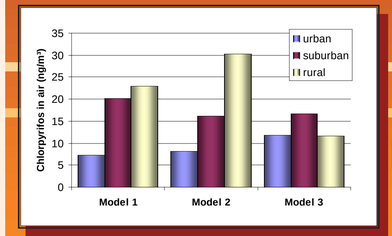

Figure 1a
Chlorpyrifos in indoor air.

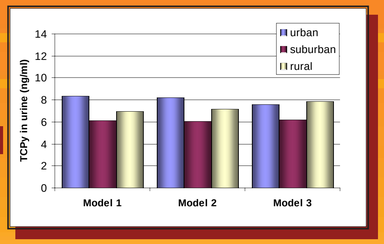

Figure 1b
TCPy in urine.

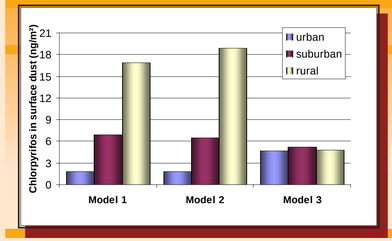

Figure 1c
Chlorpyrifos in surface dust.

Discussion

Simultaneous analysis of data from multiple groups (for example: urban, suburban, and rural dwellers) is advantageous when at least one of the variance components can be assumed to be homogeneous across the groups. Pooling data across groups to estimate common variance components requires fewer parameters to be estimated and effectively increases sample size, precision, and the power to detect effects. However, incorrectly assuming homogeneity of variance components can introduce bias and lead to erroneous conclusions.

Estimates of fixed effects (i.e., group means) and variance components may be substantially different (see Figure 1a, chlorpyrifos in air) for different covariance structures. In addition to group means, accurate estimates of the variance components themselves are often important for determining probabilities of 'exceedance' and 'overexposure', intraclass correlation coefficients of reliability, and variance ratios. These estimates can vary widely with different covariance structures. Thus, it is important to choose the covariance structure that provides the best fit to the data.

Indices of relative goodness-of-fit, such as AIC the SBC, are often useful for choosing among structures but cannot indicate whether two competing structures are significantly different. A more rigorous test, the Likelihood Ratio Test (Tables 3a-c), can be applied to conclusively determine the best fit. The test is straightforward and easy to perform.

These results indicate that the choice of covariance structure can have a large effect on the conclusions drawn from an analysis and must be given appropriate consideration using appropriate tools.

References and Resources

Cnaan A, Laird NM, Slasor P. Tutorial in biostatistics: Using the general linear mixed model to analyse unbalanced repeated measures and longitudinal data. Statist Med. 16:2349-2380 (1997).

Khattree R, Naik DN. Applied Multivariate Statistics with SAS Software, second edition, Cary, NC: SAS Institute Inc., 330 pp. (1999).

Kowalchuk RK, Keselman HJ. Mixed-model pairwise multiple comparisons of repeated measures means. Psychological Methods. 6(3):282-296 (2001).

Littell RC, Pendergast J, Natarajan R. Tutorial in biostatistics: Modelling covariance structure in the analysis of repeated measures data. Statist Med. 19:1793-1819 (2000).

Littell RC, Milliken GA, Stroup WW, Wolfinger RD. SAS System for Mixed Models, Cary, NC: SAS Institute Inc., 633 pp. (1996).

MacIntosh DL, Needham LL, Hammerstrom KA, Ryan PB. A longitudinal investigations of selected pesticide metabolites in urine. J Expos Anal Environ Epidemiol. 9(5):494-501, 1999.

Rappaport SM, Weaver M, Taylor D, Kupper L, Susi P. Application of mixed models to assess exposures monitored by construction workers during hot processes. Ann occup Hyg. 43 (7):457-469 (1999).

Strom DJ, Stansbury PS. Determining parameters of lognormal distributions from minimal information. Am Indust Hyg Assoc J. 61:877-880 (2000).

Weaver MA, Kuppper LL, Taylor D, Kromhout H, Susi P, Rappaport SM. Simultaneous assessment of occupational exposures from multiple worker groups. Ann occup Hyg. 45 (7):525-542 (2001).

Wolfinger RD. An example of using mixed models and proc mixed for longitudinal data. J Biopharmaceutical Stat. 7(4):481-500 (1997).

NOTICE
This work has been funded wholly or in part by the United States Environmental Protection Agency. It has been subjected to Agency review and approved for publication.